

AI.BACKBONE

How the harness works

A closer look at the mechanics: why a harness beats a bigger agent, how orchestrator, workflows, stations and agents fit together, what a skill is, and which workflows are available today.

30

workflows

34

agents

197

skills

91

prompts


34

tools

1 A harness is not a bigger agent


Why constraint is what makes AI usable at scale.

Give a capable model a vague goal and every run becomes a new negotiation: different steps, different outputs, different risks. A harness works like a production line. The same stations run in the same order, produce the same artefacts and stop at the same checkpoints, whoever launches it.



An unconstrained agent

every run a new negotiation




A harness

same stations, same artefacts, same checkpoints

Constraint is what makes it usable by thousands of people at once.

2 How the pieces fit together

Orchestrator, workflow, station, agent: each layer constrains the next.




-  **Orchestrator**
chooses the workflow
-  **Workflow**
orders the stations
-  **Station**
one step, its skills, one gate
-  **Agent**
model + tools, applying the station's skills

Each layer only exists to constrain the one inside it.

3 Anatomy of a station

Six fields define every step, starting with its agent and its skills.


A station is one step of a workflow. Its agent does the work by applying the skills the station names, and a gate at the exit checks the result. Nothing about it is left implicit.



4 Skills: what each station knows

What each station knows: the task at hand, described in detail.

Every station names the skills its agent must apply. A skill is a very detailed description of one engineering task: what to read, what to achieve, what never to do, which steps to follow and how the result is checked. It frames the work of the language model so it executes our method instead of improvising its own.



- **Loaded station by station.** The agent receives only the skills for the step at hand, and the same skill is reused by every station and workflow that needs it.
- **Know-how as code.** Skills are plain text files, versioned and reviewed like source code.
- **Expertise that travels.** What our specialists know about security, accessibility or data protection reaches every project, and every agent.
- **Smaller models, same quality.** The more precise the skill, the less the model has to infer, so cheaper models can do the job.
- **Adaptable to you.** Skills are tailored to each client's standards, architecture and sector rules.

Agent

=

LLM

+

tools

+

skills

skill · unit-test-generation

ILLUSTRATIVE EXCERPT

INPUTS

Approved specification, source module, test conventions

MISSION

Cover every acceptance criterion with a unit test

GUARDRAILS

Never change production code. No real personal data in fixtures

STEPS

Map criteria → write tests → run → fix → report

QUALITY CONTROLS

All tests pass, coverage threshold met, each test traced to a criterion

5 Explore the workflows

Assessment, delivery and modernisation workflows, ready to launch.

Each workflow is a sequence of stations that guides a project from idea to delivery. The orchestrator selects the right one for the task and runs it step by step. A selection:

ASSESSMENT

Brownfield Analysis
Discover repositories, index code, convert documents and reconstruct functional and architectural documentation from an existing system.
9 stations



Maturity Assessment
Assess SDLC maturity across dimensions and produce a scored report with an improvement roadmap.
4 stations



DELIVERY

BMAD
Breakthrough Method for Agile AI-Driven Development: specialised agents take a project from brief to architecture and story-by-story delivery.
4 stations



Bug Fixing
From triage through root-cause analysis, fix, regression testing and quality validation.
7 stations



Feature Implementation
End-to-end feature delivery: constitution, specification, planning, implementation and quality validation.
11 stations



Idea to Spec
Turn a raw idea into a validated, unambiguous specification with NFRs and an architecture sketch.
7 stations



Implementation Loop
Agent-assisted development loop from task selection through code generation, review, testing and commit readiness.
6 stations



Incident Resolution
Structured incident diagnosis and resolution, through fix, regression testing and knowledge capture.
7 stations



Spec Kit
Specification-only flow producing a complete spec package, for brownfield and greenfield.
9 stations



Spec to Execution
Turn a validated specification into an executable plan with risk analysis, rollback strategy and decomposed tasks.
6 stations



MODERNISATION

Modernisation
From baseline assessment through target definition, migration planning, implementation and quality assurance.
11 stations



6 Every promise is a property

Six guarantees, enforced by the engine.

Not values on a poster. Fields, checked by the engine.



Predictable

stations in a fixed order



Reliable

gates that block, not warn



Traceable

outputs declared up front



Controllable

a checkpoint at every station



Safe

an explicit tool allowlist



Repeatable

same artefacts, any team